A mean field model for an optical switch with a large number of wavelengths and centralized partial conversion

Available online at www.sciencedirect.com
ScienceDirect
Performance Evaluation 00 (2010) 1–21
Performance Evaluation

A Mean Field Model for an Optical Switch with a Large Number of Wavelengths and Centralized Partial Conversion☆

Juan F. Pérez, Benny Van Houdt

Performance Analysis of Telecommunication Systems, Department of Mathematics and Computer Science, University of Antwerp - IBBT, Middelheimlaan 1, B-2020 Antwerp, Belgium

Abstract

This paper analyzes an optical switch with centralized partial wavelength conversion by means of a mean field model. The model can be used to approximate the behavior of a switch with a large number of output wavelengths, and it becomes more accurate as the number of wavelengths increases. At each wavelength packets arrive according to a Markovian arrival process, and their size follows a general distribution with finite support. Moreover, these traffic characteristics may be different for each output port. The model provides insight into the effect of the traffic parameters on the packet loss probability, which is considered the main performance measure. In particular, we have found that, if the arrival process is Bernoulli, the loss probability is affected by the packet-size distribution only through its mean. This is no longer the case if the arrivals follow a more general Markovian process, although we have found that even in this case the loss probability is hardly sensitive to the packet-size distribution. Also, under Bernoulli arrivals we provide a closed expression for the minimum conversion ratio required to attain zero losses when the number of wavelengths tends to infinity. For Markovian arrivals we are able to compute this ratio with a single run of the mean field model.

Keywords: Optical Switch, Wavelength Conversion, Mean Field Model
2010 MSC: 68M10, 60K35

1. Introduction

Wavelength division multiplexing (WDM) has enabled optical fibers to carry huge amounts of traffic by allowing a single fiber to transmit several signals simultaneously using different

☆This work has been supported by the FWO-Flanders through project "Stochastic modeling of optical buffers and switching systems based on Fiber Delay Lines" (G.0538.07).

Email addresses: juanfernando.perez@ua.ac.be (Juan F. Pérez), benny.vanhoudt@ua.ac.be (Benny Van Houdt)

wavelengths. In addition, the transmission speed per signal keeps on rising, already achieving 40 Gb/s. To keep up with this increasing speed, the switches in the backbone network need to minimize the opto-electronic translations, as they add unnecessary delays. To this end optical switching technologies, such as OPS (optical packet switching) and OBS (optical burst switching), have been proposed. In OPS the packets are completely processed in the optical domain, while in OBS this is done for the payload only (the header is processed electronically). In both cases, the opto-electronic translations are reduced as the payload is switched in the optical domain, where contention may arise and must be resolved. If two packets require transmission through the same output port using the same wavelength, there are three main alternatives to avoid dropping one of them: one of the packets can be sent through a different output port (deflection routing), one can be buffered (optical buffering), or one can be converted to a different (idle) wavelength in the same output port (wavelength conversion). Each of these alternatives has its own advantages and disadvantages [1], and the focus of this paper will be on the use of wavelength converters alone to resolve contention. This appears to be the most accepted solution given the lack of optical random access memory and the shortcomings of deflection routing (additional load in the network, unordered arrival of packets at the destination nodes, and extra delays).

Three main attributes can be used to characterize the converters in a switch: their location; their range; and their number. The converters can be located in a centralized pool, where packets from all the ports can be translated; or they can be split in separated pools, one per port, in a non-centralized fashion. The centralized design has the potential benefit of reducing the total number of converters required in the switch, although it implies a more complex switching matrix. With regard to the conversion range, the converters may provide *full-range* conversion, i.e., a packet can be translated from and to any wavelength, or *limited-range* conversion, where a packet can only be translated to a small set of wavelengths. Finally, a switch can be designed to have as many converters as output wavelengths, a case referred to as *full* conversion (not to confuse with full-range). However, an economically feasible solution should consider fewer converters than wavelengths, in which case the switch is said to provide *partial* conversion. In this paper we introduce a model for centralized full-range partial wavelength conversion. In fact, the non-centralized architecture can be analyzed as a special case within our model.

Related Work:. One of the first analytical models to evaluate the effect of centralized partial wavelength conversion was proposed in [1], where the switch has a slotted (synchronous) operation and the packet size is fixed and equal to the slot length. Also, the authors assume that only one wavelength is used for transmission, while the others are used for contention resolution alone; therefore the load per wavelength is very small. An exact analysis for the asynchronous case was first presented in [2], where the switch is assumed to provide non-centralized partial conversion, and the wavelengths are fed by a single general arrival process. Based on this model, an approximation was introduced in [3] for the case of centralized partial conversion, which is shown to provide good results when the number of ports is large. In [2, 3] the packet length is limited to an exponential distribution, and the analysis is purely numerical. As will be described later, we will not only provide a numerical procedure to compute the loss probability under general arrival processes, but also provide a closed-form expression for this performance measure under Bernoulli arrivals. As the number of wavelengths that can be carried by a single fiber has increased significantly, the case of switches with a large number of wavelengths per port has also received attention. This case is treated in [4], where a queueing network model is proposed for a switch with non-centralized full wavelength conversion. In [4], as well as in [2], it was pointed

out that the packet loss probability was hardly sensitive to the packet-size distribution when the switch has a large number of wavelengths. Here we will provide further support to this property, as the closed-form expression for the loss probability, which holds under Bernoulli arrivals, depends on the packet-size distribution only through its mean. Many other (exact and approximate) models have been proposed to analyze optical switches with other contention resolution capabilities. For instance, the analysis of limited-range conversion has been mostly limited to approximations or simulation models due to the complex interaction of adjacent wavelengths. Also, optical buffering has been mostly considered in the absence of converters, due to the dimensionality problems arising when both solutions are combined. Finally, deflection routing has been analyzed with simulation models as its performance must be measured at the network level, which is known to be poor at high loads [5]. We do not consider these models any further as they are scarcely related to the architecture considered here.

Mean field models have been extensively used in statistical physics [6, 7] to describe the behavior of gases and other systems of particles where the interaction among these is assumed to be weak. In recent years, there has been an increasing interest in using mean field models to analyze systems such as buffers implementing active queue management with multiple TCP connections [8], networks of queues with load balancing mechanisms [7], networks with many classes of TCP-like permanent connections [9], medium access control protocols [10], reputation systems in ad-hoc networks [11], among others.

Our Contribution:. In this paper we introduce a mean field model for an optical switch with centralized partial wavelength conversion. The main feature of the mean field model is that it is exact when the number of wavelengths is infinite, and it can be used to approximate the performance of a switch with a large number of wavelengths. Modeling a switch with multiple ports and many wavelengths per port presents several difficulties, particularly with regard to the multidimensional nature of the model, as it requires keeping track of the state of the wavelengths in each port and the converters in the centralized pool. These dimensionality problems caused by the number of wavelengths are avoided by using the mean field model, and in fact the model becomes more accurate as the number of wavelengths increases. The switch is assumed to work in a synchronous manner, where the time is divided in equally-spaced slots. At each wavelength packets arrive according to a Markovian arrival process, and their size follows a general distribution with finite support. Moreover, these traffic characteristics may be different for each output port, a scenario that we refer to as heterogeneous traffic. As stated above, the non-centralized architecture can also be analyzed with our model. The assumption of a general packet-size distribution, instead of fixed size, has the advantage of reducing header processing and being more suitable for IP traffic [12].

With this model we have been able to assess how the traffic parameters influence the packet loss probability, which is considered the main performance measure. The most relevant insights gathered in this direction are: (i) Under Bernoulli arrivals, we have found an expression for the loss probability, which shows that it is affected by the packet-size distribution only through its mean. Although this is no longer the case if the arrivals follow a more general Markovian process, we have found experimentally that even in this case the loss probability is hardly sensitive to the packet-size distribution. This confirms previous observations in the same direction[4, 2]. (ii) We also provide a closed-form expression for the minimum conversion ratio required to attain zero losses when the number of wavelengths tends to infinity, under Bernoulli arrivals. This minimum conversion ratio is shown to depend on the (squared) load and the mean packet size. (iii) When the number of converters is under-dimensioned, becoming the main cause of packet losses, the

system shows a periodic behavior, where the period is equal to the greatest common divisor of the sizes of all the packets entering the switch. (iv) The burstiness of the arrival process appears to have a major effect on the loss probability, especially for mid loads. (v) When the number of wavelengths tends to infinity and the traffic among the ports is homogeneous, there is no difference between the performance of the centralized and non-centralized architectures. However, for a fixed finite number of wavelengths per output port, the results of our model are more accurate in the centralized case. (vi) Under heterogeneous traffic conditions (packet-size distribution and burstiness), important gains in conversion resources can be obtained by using a centralized architecture, even for a large number of wavelengths.

The remainder of the paper is organized as follows. Section 2 describes in detail the architecture of the switch, while Section 3 introduces the mean field model under some simplifying assumptions. The final section is concerned with numerical results that illustrate the behavior of the mean field model and uses it to analyze the effect of various traffic parameters on the switch performance.

2. The switch architecture

This section describes the architecture and operation of the optical switch under analysis. The switch, shown in Figure 1, consists of K input/output ports, each connecting to a fiber carrying W wavelengths, and a centralized pool of C converters. An incoming packet will attempt transmission through an specific output port using the same wavelength in which it entered the switch. If that wavelength is busy the packet will be converted to an available wavelength in the same output port. This conversion is performed by an idle converter in the shared pool. If all the converters or all the other wavelengths in the same output port are busy, the packet must be dropped. The probability that a packet is dropped is considered the main performance measure of the switch. With regard to the converters, we assume that they provide full-range conversion and their number is defined as a proportion of the total number of output wavelengths, i.e., $C = \sigma KW$, where σ is called the *conversion ratio*. When $\sigma = 1$ the system is said to provide *full* conversion. As stated before, an economically-feasible solution for optical switching including converters should limit the use of these devices while guaranteeing a minimal packet loss probability. Therefore our focus here is on the *partial* conversion case, where $0 < \sigma < 1$.

The switch has a synchronous operation with equally-spaced time slots and the state of the switch is observed at slot boundaries. This type of operation implies a simpler switching matrix compared to the asynchronous case, but it requires packet synchronization and alignment [13, 2]. Each wavelength in output port k has its own arrival process, modeled as a Discrete Markovian Arrival Process (DMAP) [14] characterized by the $m_k \times m_k$ matrices $D_0^{(k)}$ and $D_1^{(k)}$, for $k = 1, \ldots, K$. A DMAP$(D_0^{(k)}, D_1^{(k)})$ is a versatile point process driven by an underlying Markov chain with transition matrix $D^{(k)} = D_0^{(k)} + D_1^{(k)}$. The entries of the matrices $D_0^{(k)}$ and $D_1^{(k)}$ hold the transition probabilities of the underlying chain associated with zero and one arrivals, respectively. This kind of arrival process is computationally tractable as well as versatile, including widely used arrival processes, such as the Bernoulli process and the discrete-time versions of the interrupted Poisson process (IPP) and the Markov modulated Poisson process (MMPP), as special cases. In particular, it has been previously used to model the behavior of a bufferless asynchronous optical switch with non-centralized converters [2, 4, 15]. The size of the packets directed to output port k follow a general distribution with finite support Ξ_k. Let $r_i^{(k)}$ be the probability that a packet for output port k is of size i, for $i \in \Xi_k$, and let $L_{\max}^{(k)}$ be the maximum packet

size. Therefore, both the arrival process and the packet-size distribution may differ among output ports.

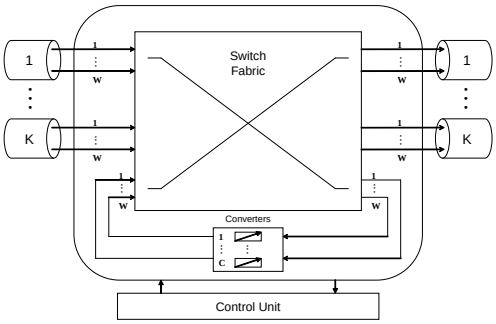

Figure 1: Optical switch with K input/output ports, W wavelengths and shared C converters

To describe the state of a wavelength or a converter we employ the scheduling horizon. This is defined as the time (number of slots) required by the wavelength/converter to transmit/translate the packet it is currently busy with. Therefore, if a wavelength has a horizon equal to 0 (i.e., it is available) and a new packet of size j arrives, the wavelength will accept the packet and update its horizon to j. Then, the horizon will be reduced by one at every slot until it reaches zero again. During the time the wavelength has a horizon greater than zero, new packets may attempt transmission using the same wavelength. If this occurs, those packets must be converted to another wavelength using an available converter, i.e., a converter with horizon equal to zero. If there is both an idle wavelength in the same output port and an available converter, these two resources are seized by the packet and their horizons are both updated from zero to the value of the packet size. The model to be introduced in the next section relies on this description to represent the state of the switch.

3. The Basic Model

In this section we introduce a simplified version of the model for the optical switch. We consider the special case where all the output ports have the same inter-arrival time (IAT) and packet-size distributions (the subscript k will be removed from the parameters of these distributions). Also, we assume that the IATs follow a geometric distribution with parameter p, and therefore the arrival process at each wavelength is a Bernoulli process. We start by describing the model for a finite number of wavelengths and then consider the limit when this number tends to infinity. Finally, we describe some results concerning a fixed point of the model.

In this model, each of the KW output wavelengths in the switch has its own Bernoulli arrival process with parameter p given by $p = \frac{\rho}{E[L]}$, where ρ is the load of the wavelength and L is the random variable describing the packet-size, with expected value $E[L]$. Each wavelength and converter will be treated as a separate object in a system of $N = KW + C$ interacting objects. To describe their evolution we observe the system at slot boundaries and consider three main steps within a slot: transmission, arrival and reallocation. During packet transmission each wavelength (resp. converter) holding a packet transmits (resp. translates) a part of it proportional to the slot length. After transmission, each wavelength may receive a new arrival with probability p. Those

packets that arrive at a busy wavelength form the set of extra-packets that must be reallocated among the idle wavelengths in the same port. While packet arrival and transmission are independent operations for each wavelength in each port, the reallocation depends on the availability of wavelengths in the same output port and converters in the shared converter pool.

To describe the evolution of the wavelengths and the converters during a time slot we consider the three steps separately and associate a transition matrix with each of them: transmission (S_l), arrivals (A_l) and reallocation (Q_l). The subscript l is equal to w or c if the matrix is associated to the evolution of a wavelength or a converter, respectively. Recall that all the output ports have the same traffic pattern (IAT and packet-size distribution) and therefore the matrices S_w and A_w describe the evolution of a wavelength in *any* of the ports. On the other hand, the evolution of a wavelength in port k during the reallocation step depends on the state of all the wavelengths in this port. Therefore, for this step we add a superscript k to the transition matrix (Q_w^k) of a wavelength in output port k. We now define these matrices explicitly and show how they are combined to describe the whole system.

S1 - Transmission:. Before transmission each wavelength/converter has a horizon between 0 and $L_{\max}$. During the transmission step (S1) the horizon of each wavelength and each converter is reduced by one. Therefore the transition matrices for wavelengths and converters are given by $S_w = S_c = T_{L_{\max}}$, where T_n is the $(n+1) \times n$ matrix with entries

$$[T_n]_{ij} = \begin{cases} 1, & i = j = 0, \\ 1, & j = i-1, i = 1, \ldots, n, \\ 0, & \text{otherwise.} \end{cases}$$

Notice that we label the rows and columns of an $m \times n$ matrix from 0 to $m-1$ and from 0 to $n-1$, respectively.

S2 - Arrivals:. When a packet arrives at an idle wavelength (horizon equal to 0) it is accepted for transmission in the same wavelength. Therefore, the first row of the matrix A_w is given by

$$[A_w]_{0j} = \begin{cases} 1-p, & j = 0, \\ pr_j, & j = 1, \ldots, L_{\max}. \end{cases}$$

If the wavelength is busy (horizon equal to $i > 0$) and there is no arrival, the horizon stays unaltered. Otherwise, the new state is $L_{\max} + i$, showing that the wavelength has scheduling horizon equal to i and holds an extra-packet for reallocation. Hence, for $i = 1, \ldots, L_{\max} - 1$,

$$[A_w]_{ij} = \begin{cases} 1-p, & j = i, \\ p, & j = L_{\max} + i. \end{cases}$$

Therefore, the size of the matrix A_w is $L_{\max} \times 2L_{\max}$. As the arrivals do not affect the state of the converters, the matrix A_c is equal to the identity matrix $I_{L_{\max}}$.

S3 - Reallocation:. The last step consists of the reallocation of the extra-packets in the available wavelengths. As mentioned before, to determine the evolution of a single wavelength or converter during this step we need to know the state of the whole system. That is, the transition matrices Q_w^k and Q_c depend on the state of the system at time t (after S2), which is described by the occupancy vector $M^N(t)$. The entries of this vector hold the proportion of the $N = KW + C$ objects that are in each possible state. For port k, let $W_k^N(t)$ be the $1 \times 2L_{\max}$ vector with entries

$w_i^k(t)$. For $i = 0, \ldots, L_{\max}$, $w_i^k(t)$ holds the number of wavelengths in port k with horizon equal to i after S2, holding *zero* extra-packets. For $i = L_{\max} + j$ and $j = 1, \ldots, L_{\max} - 1$, $w_i^k(t)$ holds the number of wavelengths in port k with horizon equal to j after S1 that received an extra-packet during S2. Similarly, let $C^N(t)$ be the $1 \times L_{\max}$ vector with entries $c_i(t)$, which hold the number of converters with horizon equal to i after S2, for $i = 0, 1, \ldots, L_{\max} - 1$. Before using the vectors $W_k^N(t)$ and $C^N(t)$ to construct $M^N(t)$ we need to normalize them by the total number of objects $KW + C$. The fraction of objects that are wavelengths in port k is $\frac{W}{KW+C} = \frac{1}{K(1+\sigma)}$, for $k = 1, \ldots, K$. Similarly, the fraction of objects that are converters is $\frac{C}{KW+C} = \frac{\sigma}{(1+\sigma)}$. We can now define the state vector $M^N(t)$ as

$$M^N(t) = \frac{1}{1+\sigma}\left[\frac{1}{K}W_1^N(t), \ldots, \frac{1}{K}W_K^N(t), \sigma C^N(t)\right],$$

which holds the proportion of objects (wavelengths and converters) in each of the $(2K + 1)L_{\max}$ possible states.

Relying on the occupancy vector $M^N(t)$, we are able to specify the entries of the matrices Q_w^k and Q_c as follows. Let $q_i^{(w,k)}(M^N(t))$ be the probability that an idle wavelength belonging to output port k receives an extra-packet of size i during reallocation at time t, for $i = 1, \ldots, L_{\max}$, and $q_0^{(w,k)}(M^N(t))$ the probability that it remains idle. Since any extra-packet is either converted to an available wavelength or dropped, it is enough to use the scheduling horizon to describe the state of a wavelength *after* this step. Therefore, the $2L_{\max} \times L_{\max}$ matrix $Q_w^k(M^N(t))$ is given by

$$Q_w^k(M^N(t)) = \left[\begin{array}{c|c|c|ccc|c}
q_0^{(w,k)}(M^N(t)) & q_1^{(w,k)}(M^N(t)) & q_2^{(w,k)}(M^N(t)) & \ldots & q_{L_{\max}-1}^{(w,k)}(M^N(t)) & q_{L_{\max}}^{(w,k)}(M^N(t)) \\
\hline
0 & 1 & 0 & \ldots & 0 & 0 \\
0 & 0 & 1 & \ldots & 0 & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & \ldots & 1 & 0 \\
0 & 0 & 0 & \ldots & 0 & 1 \\
\hline
0 & 1 & 0 & \ldots & 0 & 0 \\
0 & 0 & 1 & \ldots & 0 & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & \ldots & 1 & 0
\end{array}\right]. \quad (1)$$

Likewise, let $q_j^c(M^N(t))$ be the probability that an idle converter receives an extra-packet of size j for conversion during reallocation at time t, for $j = 1, \ldots, L_{\max}$. Then the $L_{\max} \times (L_{\max} + 1)$ matrix $Q_c(M^N(t))$ has a first row given by $[Q_c(M^N(t))]_{0j} = q_j^c(M^N(t))$, for $j = 0, 1, \ldots, L_{\max}$, where $q_0^c(t)$ is the probability that an idle converter remains idle after S3. The rest of the converters keep the same horizon, hence $[Q_c(M^N(t))]_{ii} = 1$, for $i = 1, \ldots, L_{\max} - 1$. Recall that before this step the maximum horizon a converter can have is $L_{\max} - 1$.

To determine the evolution of the idle wavelengths and converters, three quantities are relevant: the number of available converters $c_0(t)$; the number of available wavelengths $w_0^k(t)$ in output port k; and the number of extra-packets $d_k(M^N(t))$ in output port k, which is given by $d_k(M^N(t)) = \sum_{j=1}^{L_{\max}-1} w_{L_{\max}+j}^k(t)$. Since at most $w_0^k(t)$ extra-packets can be received at port k after conversion, the number of extra-packets from this port that are sent for conversion at the centralized pool is $f_k(M^N(t)) = \min\{w_0^k(t), d_k(M^N(t))\}$, and the total number of extra-packets sent for conversion is $f(M^N(t)) = \sum_{k=1}^K f_k(M^N(t))$. Thus, the number of extra-packets that is actually

converted is $g(M^N(t)) = \min\{c_0(t), f(M^N(t))\}$. Since there are no priorities among the ports, the probability that a converter that will be used by an extra-packet is assigned to an extra-packet from output port k is $f_k(M^N(t))/f(M^N(t))$. Therefore the average number of extra-packets of output port k that are converted is

$$g_k(M^N(t)) = \frac{f_k(M^N(t))}{f(M^N(t))} g(M^N(t)),$$

if $f(M^N(t)) > 0$. Otherwise, it is equal to 0. Now we can determine the time-dependent transition probabilities for an idle wavelength in output port k during S3 as

$$q_j^{(w,k)}(M^N(t)) = \begin{cases} 1 - \frac{g_k(M^N(t))}{w_0^k(t)}, & j = 0, \\ \frac{g_k(M^N(t))}{w_0^k(t)} r_j, & j = 1, \ldots, L_{\max}, \end{cases} \tag{2}$$

if $w_0^k(t) > 0$. The case where $w_0^k(t) = 0$ will be discussed in the next section. In a similar manner we can define the transition probabilities for an idle converter in this step,

$$q_j^c(M^N(t)) = \begin{cases} 1 - \frac{g(M^N(t))}{c_0(t)}, & j = 0, \\ \frac{g(M^N(t))}{c_0(t)} r_j, & j = 1, \ldots, L_{\max}, \end{cases} \tag{3}$$

if $c_0^k(t) > 0$. As with the idle wavelengths, the case $c_0^k(t) = 0$ will be treated in the next section. The factor r_j comes from the fact that the $g(M^N(t))$ converted extra-packets are selected randomly among the set of extra-packets, therefore keeping the same packet-size distribution. This concludes the definition of the transition matrices for this step. We now combine the three steps to describe the evolution of the system during a slot.

3.1. Combining S1, S2 and S3 - Mean Field

By observing the system after S2, we define the transition matrix at time t for a single wavelength in output port k as $R^N_{(w,k)}(M^N(t)) = Q^k_w(M^N(t))S_wA_w$, and for a single converter as $R^N_c(M^N(t)) = Q_c(M^N(t))S_cA_c$. These $K + 1$ matrices can be combined into a single matrix describing the evolution of a single object at time t,

$$R^N(M^N(t)) = \begin{bmatrix} R^N_{(w,1)}(M^N(t)) & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & R^N_{(w,K)}(M^N(t)) & 0 \\ 0 & \cdots & 0 & R^N_c(M^N(t)) \end{bmatrix}. \tag{4}$$

From their the definition, it is clear that the entries of S_w, A_w, S_c, A_c, as well as all but the first-row entries of Q^k_w and Q_c, are independent of the number of objects (wavelengths and converters) in the system. Actually, the first-row entries of Q_w and Q_c are also independent of the number of objects. To see this, let $\vec{m}$ be a $1 \times (2K + 1)L_{\max}$ occupancy vector partitioned as

$$\vec{m} = [\vec{m}^w_1, \vec{m}^w_2, \ldots, \vec{m}^w_K, \vec{m}^c].$$

Since the proportion of objects that are wavelengths in output port k is $K(1 + \sigma)$, the vector $K(1 + \sigma)\vec{m}^w_k$ is a $1 \times 2L_{\max}$ occupancy vector. For $i = 0, \ldots, L_{\max}$, the i-th entry of this vector

holds the proportion of wavelengths in output port k with horizon equal i, holding zero extra packets, before reallocation. For $i = L_{\max} + j$, and $j = 1, \ldots, L_{\max} - 1$, the i-th entry of the same vector holds the proportion of wavelengths in output port k with horizon equal to j before reallocation, holding an extra packet. Similarly, the vector $\frac{1+\sigma}{\sigma}\vec{m}^c$ is a $1 \times L_{\max}$ occupancy vector, the i-th entry of which holds the proportion of converters in the centralized buffer with horizon equal to i before the reallocation step, for $i = 0, \ldots, L_{\max} - 1$. Now, we can redefine the entries of Q_w in terms of the vector $\vec{m}$ as

$$q_j^{(w,k)}(\vec{m}) = \begin{cases} 1 - \frac{g_k(\vec{m})}{[\vec{m}_k^w]_0}, & j = 0, \\ \frac{g_k(\vec{m})}{[\vec{m}_k^w]_0} r_j, & j = 1, \ldots, L_{\max}, \end{cases} \tag{5}$$

for $[\vec{m}_k^w]_0 > 0$. As in the previous section, to define $g_k(\vec{m})$ we first need to introduce $d_k(\vec{m}) = \sum_{j=1}^{L_{\max}-1} [\vec{m}_k^w]_{L_{\max}+j}$, and $f_k(\vec{m}) = \min\{[\vec{m}_k^w]_0, d_k(\vec{m})\}$. Also, let $f(\vec{m}) = \sum_{k=1}^{K} f_k(\vec{m})$, and $g(\vec{m}) = \min\{[\vec{m}^c]_0, f(\vec{m})\}$. Now, we can define $g_k(\vec{m})$ as $g_k(\vec{m}) = \frac{f_k(\vec{m})}{f(\vec{m})} g(\vec{m})$, if $f(\vec{m}) > 0$, and 0 otherwise. Clearly, the definition of the entries of $Q_w(\vec{m})$ can be obtained from Eq. (2) by simply dividing by N all the quantities involved in their computation. Moreover, the entries of $Q_w(\vec{m})$ are independent of the number of objects, as they can be defined completely as a function of the proportion (not the number) of objects in each possible state. Since a similar construction can be made for the entries of $Q_c(\cdot)$, we conclude that the entries of the matrix $R^N(\cdot)$ are also independent of the number of objects. We therefore remove the superscript N, and refer to this matrix simply as $R(\cdot)$.

This system, composed of wavelengths and converters, can be analyzed under the framework introduced in [16] for a general system of interacting objects. When the number of objects tends to infinity and under some mild conditions, such a system is shown to converge to its mean field [16], which is a time-dependent deterministic model. In our case, the objects are of $K + 1$ different classes, their status at time t is contained in the vector $M^N(t)$ and the evolution of the system is described by the matrix $R(\vec{m})$. The main convergence result in [16] is re-stated here, in a simplified form that suffices for the model at hand, as Theorem 1, which relies on the following condition.

Condition 1 (Hypothesis **H** in [16])**.** *For all i, j, as $N \to \infty$, $[R^N(\vec{m})]_{ij}$ converges uniformly in $\vec{m}$ to some $[R(\vec{m})]_{ij}$, which is a continuous function of $\vec{m}$.*

Theorem 1 (Theorem 4.1 in [16])**.** *Assume that the initial occupancy measure $M^N(0)$ converges almost surely to a deterministic limit $\mu(0)$. Define $\mu(t)$ iteratively from its initial value $\mu(0)$, for $t \geq 0$, as $\mu(t+1) = \mu(t)R(\mu(t))$. Then, for any fixed time t, almost surely, $\lim_{N\to\infty} M^N(t) = \mu(t)$.*

Therefore, to apply Theorem 1 we need to ensure the almost sure convergence of $M^N(0)$ to $\mu(0)$, as $N \to \infty$, as well as to verify that Condition 1 holds. The convergence of $M^N(0)$ is imposed by simply assuming that the system always starts in an empty state, as follows. Let $e_1^{L_{\max}}$ be the $1 \times L_{\max}$ vector with 1 in the first position and zero everywhere else. Let the initial state of the system be given by $C^N(0) = e_1^{L_{\max}} A_c$ and $W_k^N(0) = e_1^{L_{\max}} A_w$, $k = 1, \ldots, K$, which corresponds to an empty system and is independent of N. We now set $\mu(0) = M^N(0)$ for every $N > 0$, and the convergence is trivially satisfied. The compliance with Condition 1 is stated in the following theorem.

Theorem 2. *The entries of the matrices $R(\vec{m})$, defined in* (4)*, comply with the requirements of Condition 1, whenever $[\vec{m}_k^w]_0 > 0$, for $k = 1, \ldots, K$, and $[\vec{m}_c]_0 > 0$. Also, if any of these entries*

is equal to zero, the result in Theorem 1 also holds. Therefore, for any fixed time t, almost surely, $\lim_{N\to\infty} M^N(t) = \mu(t),\ t \geq 0$.

Proof. As we have shown, the entries of the matrix $R(\vec{m})$ are independent of the number of objects N. Therefore, the uniform convergence condition is trivially satisfied. Regarding the continuity condition, we observe that the matrices S_w, S_c, A_w and A_c are independent of $\vec{m}$, hence the entries of $R_{(w,k)}(\vec{m})$ and $R_c(\vec{m})$ are a weighted sum of the first-row entries of $Q^k_w(\vec{m})$ and $Q_c(\vec{m})$, respectively. As these entries are defined, in (2) and (3), as sums, products and ratios of the entries of $\vec{m}$, these are continuous functions of $\vec{m}$, as long as $[\vec{m}^w_k]_0 > 0$, for $k = 1, \ldots, K$, and $[\vec{m}_c]_0 > 0$. As a result, if $[\vec{m}^w_k]_0 > 0$ and $[\vec{m}_c]_0 > 0$, the entries of $R(\vec{m})$ comply with Condition 1.

For the case when either $[\vec{m}^w_k]_0$ or $[\vec{m}_c]_0$ are zero, we need to reconsider the proof of Theorem 1 in [16], and particularly Lemma 8.3 therein. To this end, let us define $X^N_n(t)$ as the state of object n at time t in a system with N objects, for $n = 1, \ldots, N$. The initial state of these objects $(X^N_n(0))$ is chosen to comply with $M^N(0)$, which, as stated before, corresponds to an empty system. Recall that the initial state of the mean field also corresponds to an empty system, hence $\mu(0) = M^N(0)$, for $N > 0$. Now let $P^N(i, \vec{m}, u)$ be a program that computes the next state of an object in a system with N objects, i.e., relying on $R^N(\vec{m})$, given that the current state of the object is i, the occupancy measure is $\vec{m}$, and a random number u is drawn uniformly from $(0, 1)$. The evolution of the objects in the system is thus given by $X^N_n(t+1) = P^N(X^N_n(t), M^N(t), U_n(t))$, with $\{U_n(t), 1 \leq n \leq N, t \geq 0\}$ i.i.d. random variables uniformly distributed in $(0, 1)$. Now, define another collection of interacting objects $\tilde{X}^N_n(t)$, such that $\tilde{X}^N_n(0) = X^N_n(0)$, for $1 \leq n \leq N$. The evolution of these objects is ruled by $P(i, \vec{m}, u)$, which is defined in a similar manner as $P^N(i, \vec{m}, u)$, but in this case the evolution does not depend on $R^N(\vec{m})$, but on $R(\vec{m})$. In our case, both programs are identical, since $R^N(\vec{m})$ is independent of N and therefore equal to $R(\vec{m})$. However, what differentiates $\tilde{X}^N_n(t)$ from $X^N_n(t)$ is that the former depends on the mean field occupancy vector $\mu(t)$, instead of the $M^N(t)$ occupancy measure. Thus, this new set of objects evolves as $\tilde{X}^N_n(t+1) = P(\tilde{X}^N_n(t), \mu(t), U_n(t))$.

In Lemma 8.3 in [16] it is proven that $\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} 1\{X^N_n(t) \neq \tilde{X}^N_n(t)\} = 0$ almost surely, where $1\{\cdot\}$ is the indicator function. This boils down to prove that $\lim_{N\to\infty} A^N_i = 0$ a.s., where A^N_i is defined as

$$A_i = \frac{1}{N} \sum_{n=1}^{N} 1\{X^N_n(t+1) \neq \tilde{X}^N_n(t+1), X^N_n(t) = \tilde{X}^N_n(t) = i\},$$

for every i in the state space. In [16], this is proven in general when the entries of $R(\vec{m})$ are a continuous function of $\vec{m}$. Recall that, in our model, this is true if $[\vec{m}^w_k]_0 > 0$, for $k = 1, \ldots, K$, and $[\vec{m}_c]_0 > 0$. However, even if $[\vec{m}^w_k]_0 = 0$ or $[\vec{m}_c]_0 = 0$ it can be shown that a.s. $\lim_{N\to\infty} A^N_i = 0$. This follows from writing

$$A_i \leq \frac{1}{N} \sum_{n=1}^{N} 1\{X^N_n(t) = \tilde{X}^N_n(t) = i\}, \tag{6}$$

and noting that if the i-th entry of either $\mu(t)$ or $M^N(t)$ is equal to zero, then the right-hand-side of (6) is also equal to zero. Therefore, also in this case $A_i \to 0$ a.s., and Lemma 8.3 in [16] follows, hence Theorem 1 holds.

□

In Theorem 2 we have shown that the entries of $R(\vec{m})$ may have a discontinuity at $[m]_i = 0$, and Theorem 1 still holds. In our case, the discontinuity arises in some points when either $[\vec{m}^w_k]_0$

or $[\vec{m}_c]_0$ are zero, and we have not provided a definition for this case. For completeness, we define $q_j^{(w,k)}(\vec{m})$, when $[\vec{m}_k^w]_0 = 0$, as

$$q_j^{(w,k)}(\vec{m}) = \begin{cases} 0, & j = 0, \\ r_j, & j = 1, \ldots, L_{\max}. \end{cases}$$

Actually, these entries could be defined differently, without affecting the dynamics of the system. Intuitively, this can be understood by observing that these entries will always be multiplied by $[\vec{m}_k^w]_0$, which is assumed to be zero, and therefore they have no effect on the evolution of the system state. A similar definition can be given for $q_j^c(\vec{m})$ when $[\vec{m}_c]_0 = 0$.

For further reference it is useful to partition the vector $\mu(t)$ according to the object classes (wavelengths and converters),

$$\mu(t) = \frac{1}{1+\sigma}\left[\frac{1}{K}\mu^{w,1}(t), \ldots, \frac{1}{K}\mu^{w,K}(t),\ \sigma\mu^c(t)\right].$$

Relying on Theorem 2, we can approximate the behavior of a switch with a large number of wavelengths (objects) by means of the mean field. However, this result is time-dependent and says nothing about the behavior of the system when t tends to infinity. This is the topic of the next section.

3.2. Combining S1, S2 and S3 - Fixed Point

We are now left with a deterministic system with initial state $\mu(0)$ and time-dependent transition matrix $R(\mu(t))$. To study the behavior of this system as a function of time we change the observation time-points. We now observe the system after S1, i.e., after transmission and just before arrivals. Let $\alpha(t)$ be the occupancy vector describing the state of all the wavelengths and converters at time t just after S1 and before S2. Similar to $\mu(t)$, the vector $\alpha(t)$ can be partitioned as

$$\alpha(t) = \frac{1}{1+\sigma}\left[\frac{1}{K}\alpha^{w,1}(t), \ldots, \frac{1}{K}\alpha^{w,K}(t),\ \sigma\alpha^c(t)\right],$$

where $\alpha^{w,k}(t)$ (resp. $\alpha^c(t)$) is an occupancy vector describing the state of the wavelengths in output port k (resp. converters in the centralized pool). Since $\alpha(t)$ and $\mu(t)$ describe the state of the switch before and after S2, respectively, $\mu(t)$ can be obtained from $\alpha(t)$ as $\mu(t) = \alpha(t)\begin{bmatrix} I_K \otimes A_w & 0 \\ 0 & A_c \end{bmatrix}$, where $\otimes$ is the Kronecker product. Now let the matrix $P_w^k(\alpha(t))$ describe the transition probabilities of the wavelengths in output port k at time t, observing the system after S1. This is given by $P_w^k(\alpha(t)) = A_w Q_w^k(\mu(t)) S_w$, where, as stated above, $\alpha(t)$ completely determines $\mu(t)$. A similar matrix $P_c(\alpha(t))$ can be specified for the evolution of the converters observed just after S1.

The main advantage of defining the $L_{\max} \times L_{\max}$ matrices $P_w^k(\alpha(t))$ and $P_c(\alpha(t))$ is that they can be expressed as

$$P_l^k(\alpha(t)) = \begin{bmatrix} p_0^{(l,k)}(\alpha(t)) & p_1^{(l,k)}(\alpha(t)) & \ldots & p_{L_{\max}-2}^{(l,k)}(\alpha(t)) & p_{L_{\max}-1}^{(l,k)}(\alpha(t)) \\ 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & 0 \end{bmatrix}, \tag{7}$$

for $l \in \{w, c\}$ and $k = 1, \ldots, K$ (if $l = c$ the sub/superscript k is obviously removed). The specific values of the first-row entries, in the case of the matrices associated to the evolution of the wavelengths, are given by

$$p_j^{(w,k)}(\alpha(t)) = \begin{cases} (1-p)(q_0^{(w,k)}(\mu(t)) + q_1^{(w,k)}(\mu(t))) + pr_1, & j = 0, \\ (1-p)q_{j+1}^{(w,k)}(\mu(t)) + pr_{j+1}, & j = 1, \ldots, L_{\max}-1. \end{cases} \tag{8}$$

For the converters those values are

$$p_j^{c}(\alpha(t)) = \begin{cases} q_0^{c}(\mu(t)) + q_1^{c}(\mu(t)), & j = 0, \\ q_{j+1}^{c}(\mu(t)), & j = 1, \ldots, L_{\max}-1. \end{cases} \tag{9}$$

These $K+1$ matrices can be recombined into a single matrix with $K+1$ irreducible classes,

$$P(\alpha(t)) = \begin{bmatrix} P_w^1(\alpha(t)) & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & P_w^K(\alpha(t)) & 0 \\ 0 & \cdots & 0 & P_c(\alpha(t)) \end{bmatrix}.$$

Let $\mathcal{M}_{(K+1)L_{\max}}$ be the set of all occupancy vectors of size $(K+1)L_{\max}$. The matrix $P(\alpha)$ defines a mapping from $\mathcal{M}_{(K+1)L_{\max}}$ into $\mathcal{M}_{(K+1)L_{\max}}$, which is said to have a fixed point α if there exists a vector $\alpha \in \mathcal{M}_{(K+1)L_{\max}}$ such that $\alpha P(\alpha) = \alpha$. To find this fixed point we set-up the system of equations $\alpha P(\alpha) = \alpha$ and solve it for α. To do so, we first need an explicit definition of the (first-row) entries of $P_w^k(\alpha)$ and $P_c(\alpha)$ in terms of α. Let the first-row entries of these matrices, as defined in Equation (7), be $p_j^{(w,k)} = p_j^{(w,k)}(\alpha)$ and $p_j^c = p_j^c(\alpha)$, for $j = 0, \ldots, L_{\max}-1$. These probabilities depend on the value of $q_j^{(w,k)} = q_j^{(w,k)}(\mu)$ and $q_j^c = q_j^c(\mu)$, for $j = 0, \ldots, L_{\max}-1$, as in equations (8) and (9). The vector μ is given by $\mu = \alpha \begin{bmatrix} I_K \otimes A_w & 0 \\ 0 & A_c \end{bmatrix}$, and therefore $\mu_0^{(w,k)} = \alpha_0^{(w,k)}(1-p)$ and $\mu_0^c = \alpha_0^c$. This equation relates the proportion of idle wavelengths and converters before and after arrivals in the fixed point. Similarly, the proportion of wavelengths in output port k holding an extra-packet after S2 in the fixed point is given by $\delta_k = (1 - \alpha_0^{(w,k)})p$.

Now let ϕ_k be the number of extra-packets in port k sent for conversion in the fixed point, as a proportion of the total number of objects, which is given by

$$\phi_k = \frac{1}{K(1+\sigma)} \min\left\{\mu_0^{(w,k)}, \delta_k\right\} = \frac{1}{K(1+\sigma)} \min\left\{\alpha_0^{(w,k)}(1-p), (1-\alpha_0^{(w,k)})p\right\},$$

and thus only depends on the vector α. Therefore, the number of converted extra-packets in the fixed point, as a proportion of the total number of objects, is given by

$$\gamma = \min\left\{\sum_{k=1}^{K} \phi_k, \frac{\sigma}{1+\sigma}\mu_0^c\right\} = \min\left\{\sum_{k=1}^{K} \phi_k, \frac{\sigma}{1+\sigma}\alpha_0^c\right\}, \tag{10}$$

which also depends on α alone. Additionally, let γ_k be the number of converted extra-packets from output port k as a fraction of the total number of objects, given by

$$\gamma_k = \frac{\phi_k}{\sum_{j=1}^{K} \phi_j}\gamma, \tag{11}$$

where the first term on the right-hand side is the probability that an extra-packet sent for conversion is actually an extra-packet from output port k. Since γ and γ_k depend only on $\alpha_0^{(w,k)}$ and α_0^c, we can use them in Equations (2) and (3) to determine the transition probabilities for the idle wavelengths in the fixed point in terms of α as

$$q_j^{(w,k)} = \begin{cases} 1 - \frac{K(1+\sigma)\gamma_k}{\mu_0^{(w,k)}} & = 1 - \frac{K(1+\sigma)\gamma_k}{\alpha_0^{(w,k)}(1-p)}, \quad j = 0, \\ \frac{K(1+\sigma)\gamma_k}{\mu_0^{(w,k)}} r_j & = \frac{K(1+\sigma)\gamma_k}{\alpha_0^{(w,k)}(1-p)} r_j, \quad j = 1, \ldots, L_{\max}, \end{cases}$$

and for the idle converters as

$$q_j^c = \begin{cases} 1 - \frac{(1+\sigma)\gamma}{\sigma\mu_0^c} & = 1 - \frac{(1+\sigma)\gamma}{\sigma\alpha_0^c}, \quad j = 0, \\ \frac{(1+\sigma)\gamma}{\sigma\mu_0^c} r_j & = \frac{(1+\sigma)\gamma}{\sigma\alpha_0^c} r_j, \quad j = 1, \ldots, L_{\max}. \end{cases}$$

With these expressions and equations (8) and (9) we find that the first-row entries of $P_w^k(\alpha)$ are

$$p_0^{(w,k)} = (1-p) - \frac{K(1+\sigma)\gamma_k}{\alpha_0^{(w,k)}}(1-r_1) + pr_1, \text{ and } p_j^{(w,k)} = \left(\frac{K(1+\sigma)\gamma_k}{\alpha_0^{(w,k)}} + p\right) r_{j+1},$$

for $j = 1, \ldots, L_{\max} - 1$. Also, the first-row entries of the matrix $P_c(\alpha)$ are

$$p_0^c = 1 - \frac{(1+\sigma)\gamma}{\sigma\alpha_0^c}(1-r_1), \text{ and } p_j^c = \frac{(1+\sigma)\gamma}{\sigma\alpha_0^c} r_{j+1}, \ j = 1, \ldots, L_{\max} - 1.$$

Now that the entries of the matrix $P(\alpha)$ are explicitly given in terms of α, γ and γ_k, we solve the system $\alpha = \alpha P(\alpha)$. Solving for $\alpha^{(w,k)}$ we find

$$\alpha_0^{(w,k)} = \frac{1 - \gamma_k(E[L]-1)K(1+\sigma)}{1-p+\rho}, \text{ and } \alpha_i^{(w,k)} = \left(\gamma_k K(1+\sigma) + p\alpha_0^{(w,k)}\right) P[L \geq i+1],$$

for $i = 1, \ldots, L_{\max} - 1$, where $P[L \geq i+1]$ is the probability that the packet size is greater than or equal to $i+1$. In a similar manner we find that the entries of the vector α^c are

$$\alpha_0^c = 1 - \frac{1+\sigma}{\sigma}\gamma(E[L]-1), \text{ and } \alpha_i^c = \frac{1+\sigma}{\sigma}\gamma P[L \geq i+1], \quad i = 1, \ldots, L_{\max} - 1.$$

We have found expressions for the fixed-point vector α but these are still in terms of γ and γ_k. We now consider each of the three possible values that γ and γ_k can take to find explicit formulas for the fixed point vector in each case.

Case 1:. The first case corresponds to a switch with an infinite number of wavelengths that has enough converters to convert every packet and is not overloaded. From equations (10) and (11) we find that

$$\gamma = \frac{p}{K(1+\sigma)} \sum_{k=1}^{K} \left(1 - \alpha_0^{(w,k)}\right), \text{ and } \gamma_k = \frac{(1-\alpha_0^{(w,k)})p}{K(1+\sigma)}.$$

Using these values we find that

$$\alpha_0^{(w,k)} = 1 - \rho + p, \text{ and } \alpha_i^{(w,k)} = pP[L \geq i+1], \ i = 1, \ldots, L_{\max} - 1.$$

And for the converters we find that

$$\alpha_0^c = 1 - \frac{1}{\sigma}\left(p(E[L]-1)\right)^2 \text{ and } \alpha_i^c = \frac{p^2}{\sigma}(E[L]-1)P[L \geq i+1], \ i = 1, \ldots, L_{\max} - 1.$$

Case 2:. The second case corresponds to a system that does not have enough converters to convert every extra-packet. In this case we find that $\gamma = \frac{\sigma\alpha_0^c}{1+\sigma}$, since the converters become the bottleneck of the system. The fixed point for the converters is given by

$$\alpha_0^c = \frac{1}{E[L]} \text{ and } \alpha_i^c = \frac{P[L \geq i+1]}{E[L]}, \quad i = 1, \ldots, L_{\max} - 1.$$

As in this case γ_k is given by $\gamma_k = \frac{1-\alpha_0^{(w,k)}}{\sum_{j=1}^{K}\left(1-\alpha_0^{(w,j)}\right)} \frac{\sigma}{1+\sigma}\alpha_0^c$, we can use the value of α_0^c to find

$$\alpha_0^{(w,k)} = \frac{E[L] - \sigma(E[L]-1)}{E[L](1-p+\rho)}, \text{ and } \alpha_i^{(w,k)} = \frac{\sigma+\rho}{E[L](1-p+\rho)} P[L \geq i+1], \; i = 1, \ldots, L_{\max} - 1.$$

Case 3:. The last case considers a heavily loaded switch where the wavelengths are not enough to handle the incoming packets, meaning an extra-packet might find a converter but not an idle wavelength in the output port. In this case we find that γ and γ_k are

$$\gamma = \frac{1-p}{K(1+\sigma)} \sum_{k=1}^{K} \alpha_0^{(w,k)} \text{ and } \gamma_k = \frac{\alpha_0^{(w,k)}(1-p)}{K(1+\sigma)}.$$

Therefore, the vector $\alpha_0^{(w,k)}$ is given by

$$\alpha_0^{(w,k)} = \frac{1}{E[L]}, \text{ and } \alpha_i^{(w,k)} = \frac{P[L \geq i+1]}{E[L]}, \; i = 1, \ldots, L_{\max} - 1.$$

And the vector α^c is given by

$$\alpha_0^c = 1 - \frac{(1-p)(E[L]-1)}{\sigma E[L]}, \text{ and } \alpha_i^c = \frac{1-p}{\sigma E[L]} P[L \geq i+1], \; i = 1, \ldots, L_{\max} - 1.$$

3.2.1. The loss probability and the optimal conversion ratio

An interesting observation is that the value of α_0^w and α_0^c does not depend on the distribution of the packet size but only on its expected value $E[L]$. This becomes relevant when looking at the loss probability of the system p_{loss}, which is the main measure of performance. The loss probability is the ratio between the average number of packets that must be dropped per time slot and the average number of packets that enter the system in each time slot. Therefore, the loss probability is given by

$$p_{\text{loss}} = \frac{\frac{\delta}{1+\sigma} - \gamma}{\frac{p}{1+\sigma}} = \frac{\delta - \gamma(1+\sigma)}{p},$$

where $\delta = \frac{1}{K}\sum_{k=1}^{K} \delta_k$. As shown above, δ and γ depend on the value of α_0^w and α_0^c alone, and therefore p_{loss} does not depend on the packet-size distribution but only on its expected value. This confirms previous observations [2, 4] related to an apparent insensitivity of the switch performance to the packet-size distribution when the number of wavelengths is large. However, this result relies on the assumption of geometrically-distributed IATs. We have observed that this result no longer holds if the IATs are described by a general DMAP, although even in this case the packet-size distribution appears to have little influence on the loss probability. This will be illustrated in the numerical results in Section 4.

Based on the previous results we can consider the question of how many converters are necessary to attain a loss probability equal to zero. Or, in others words, what is the minimum conversion ratio σ^* such that $p_{\text{loss}} = 0$. For p_{loss} to be equal to zero, γ must be equal to $\frac{\delta}{1+\sigma}$, as in case 1. If the conversion ratio is just enough to prevent any losses, then γ must also be equal to $\frac{\sigma\alpha_0^c}{1+\sigma}$, as in case 2. Therefore, we can compute the value of σ^* by equating the value of γ in these two cases. Solving this equation for σ we find that

$$\sigma^* = \rho^2\left(1 - \frac{1}{E[L]}\right). \tag{12}$$

The value of σ^* is proportional to the square of the load and if $E[L]$ tends to infinite (the slot length tends to zero) σ^* tends to ρ^2. A decrease in the mean packet-size (increase in the slot length) implies a decrease in the number of converters required to attain zero loss probability.

With a similar analysis we can consider the situation when both cases 1 and 3 occur. This is the point where the wavelengths become the bottleneck of the system. By solving a similar equation as before we find that $\rho = 1$, i.e., the system will only present losses caused by the lack of available wavelengths if it is overloaded ($\rho > 1$). This result is expected since the number of wavelengths is assumed to be infinite and, if no losses are caused by the lack of converters, the only way to observe packet losses is by having a load greater than one.

The previous discussion shows that there exists a fixed point for the mapping defined by $P(\cdot)$ and that this point can be expressed explicitly in terms of the system parameters. However, we have not shown that, starting from any vector $\alpha(0)$, the system always converges to that fixed point. In our experiments we start with an arbitrary initial state $\alpha(0)$ and let the system evolve, finding two possible behaviors. On the one hand, when the system has enough converters to prevent losses the system state converges toward a single state, which coincides with the fixed point described in this section. On the other hand, if the system presents losses due to the lack of converters, it will converge toward a set of points, which are visited periodically. Moreover, we have observed that in the latter case the period is equal to the greatest common divisor of the possible packet sizes (d), and the average of these states is equal to the fixed point described above. Therefore, in the experiments we let the system evolve until a time t such that $\|\alpha(t) - \alpha(t-d)\| < \epsilon$, with $\epsilon = 10^{-10}$, which always results in the convergence of the system state. Another important issue is to show that a sequence of finite systems in steady state with increasing number of wavelengths tends toward a limit equal to the fixed point of the mean field. We have observed through simulations that this is the case but we did not find a formal proof.

The mean field model introduced in this section can be easily generalized to consider both heterogeneous traffic, and a general DMAP arrival process per wavelength. In this more general setting, the model still falls within the framework in [16], and therefore it can be used to approximate the state of a finite system at any finite time. Moreover, we have performed a large number of experiments and found that the general model behaves similarly to the basic model introduced in the previous section, when the time tends to infinity. However, for the general model, we have not found an expression for the fixed point toward which the system state converges, in terms of the system parameters. A detailed description of the general model can be found in [17].

4. Results

In this section we start by illustrating the time-dependent behavior of the mean field model. Then we compare the stationary performance of the model against simulations of a switch with

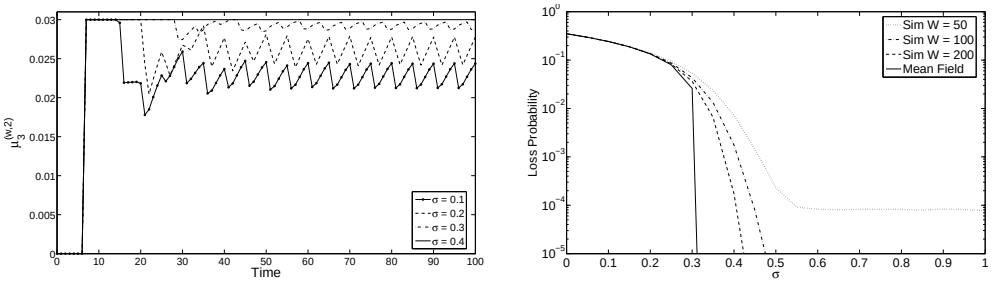

(a) $K = 2$, $L \in \{5, 15\}$, Geo IATs, $\rho = 0.6$ (b) $K = 4$, $L \sim$ Unif$(5, 15)$, Geo IATs, $\rho = 0.6$

Figure 2: Mean field and simulation results

a (large) finite number of wavelengths. Afterward we make use of the mean field model to analyze the effect of various traffic parameters on the switch performance. As mentioned before we consider the loss probability and σ^*, the minimum conversion ratio to attain a zero loss probability, as the main performance measures. We examine the effect of the packet-size distribution, the load and the arrival process' burstiness, for both the homogeneous and heterogeneous traffic cases.

4.1. Validation

We start by looking at the time-dependent behavior of the mean field model, as shown in Figure 2(a). In this scenario the switch has two output ports, homogeneous traffic, geometric IATs, load equal to 0.6, and the packet size is either 5 or 15 with equal probability. In Figure 2(a) we depict, as a function of time, the fraction of wavelengths in output port 2 with scheduling horizon equal to 3 ($\mu_3^{(w,2)}$). This selection is arbitrary as any other entry in the state vector shows a similar behavior. Since the mean packet size $E[L]$ is 10 and the load ρ is 0.6, we know from Equation (12) that the optimal conversion ratio σ^* is 0.324. Hence, if the conversion ratio is less than this value the system will present losses due to the lack of converters. In this figure we observe that when σ is below 0.324 the system tends to a periodic state, and the period is equal to 5, which is the greatest common divisor of the packet sizes (5 and 15). Also, when the conversion ratio is closer to 0.324 we see that the fluctuations are smaller, and when σ surpasses σ^* the system converges to a single fixed point. We have found the same behavior in a large number of experiments, from which we have concluded that the system may converge either to a single or to a periodic steady state. The former case occurs when the system has enough converters to prevents losses, while the latter is the result of an under-dimensioned conversion ratio. Additionally, when the latter case occurs, the number of different states that the system visits in steady state is equal to the greatest common divisor of the possible packet sizes.

We now compare the loss probability in a finite simulated system with the one computed with the mean field model. In Figure 2(b) we show this comparison for a switch with 4 ports, the packet size is uniformly distributed between 5 and 15, the IATs are geometrically distributed and the load per wavelength is 0.6. In this figure the loss probability is depicted against the conversion ratio σ. We observe how the performance of the finite systems tends to that of the mean field when the number of wavelengths per port increases, in this case from 50 to 200. If the conversion ratio is equal to one, and the traffic is uniform with geometric IATs, a finite switch behaves as a Geo/Geo/KW/KW loss queue, and therefore, as the conversion ratio increases, the loss

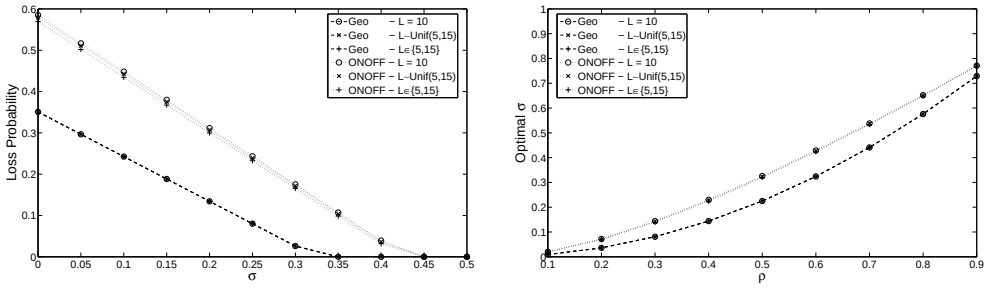

(a) Loss probability. $\rho = 0.6$, $\gamma = 5$

(b) Optimal σ. $\gamma = 5$

Figure 3: Effect of the packet-size distribution

probability of a finite switch converges toward that of the loss queue. For instance, for $W = 50$ we observe that this minimum loss probability is around 10^{-4}, and for a conversion ratio of 0.55 the loss probability of the finite system has already reached a value very close to the minimum. A similar behavior occurs for $W = 100$ and $W = 200$ but in this case the loss probability is so small that simulations become computationally prohibitive.The main difference between the mean field and a finite system is that the mean field model can be dimensioned to attain zero loss probability (σ^*), while the conversion ratio in any real finite system will be dimensioned to attain a very small loss probability ($\hat{\sigma}$). Although σ^* will typically be an optimistic value for $\hat{\sigma}$, it is a very close approximation, especially if the number of wavelengths is large. Therefore, the mean field model can be used to provide a fast-to-compute value to start the search for the actual $\hat{\sigma}$, thus restricting the search for the value of $\hat{\sigma}$ to a small neighborhood above σ^*. Another very relevant feature of the mean field model is that it let us analyze the effect of various traffic parameters on the performance of the switch, without performing any time-consuming simulations, which become more expensive as the number of wavelengths increases, particularly if the measure of interest (the loss probability) is very small. The remainder of this section is devoted to the analysis of the effect of the traffic parameters on both the loss probability and σ^*.

4.2. Homogeneous traffic

In this section we consider the case where the traffic is homogeneous, i.e., the traffic to all the output ports has the same characteristics. Figure 3 illustrates the effect of the packet-size distribution on the switch performance. As it was established in Section 3, when the IATs follow a geometric distribution, the only influence of the packet-size distribution is through its mean. In this case we change the distribution without altering the mean, and Figure 3(a) shows that this has no effect on the loss probability when the IATs are geometrically distributed. We also consider a more general arrival process, called ON-OFF, which is a particular case of a DMAP. An ON-OFF process has an underlying chain with two states: in the so-called ON state the process generates arrivals with geometric IATs, while in the OFF state no arrivals are generated. This kind of arrival process has been previously used to model the arrival process at an optical switch [15, 4]. The duration of the ON and OFF periods is geometrically distributed, with the mean duration of the OFF periods being γ times that of the ON periods. In Figure 3 we also observe that for ON-OFF arrivals with $\gamma = 5$ the effect of the packet-size distribution is rather small, even though the distributions we consider are significantly different. The three packet-size distributions are: deterministic with packet size $B = 10$; uniformly distributed between 5 and

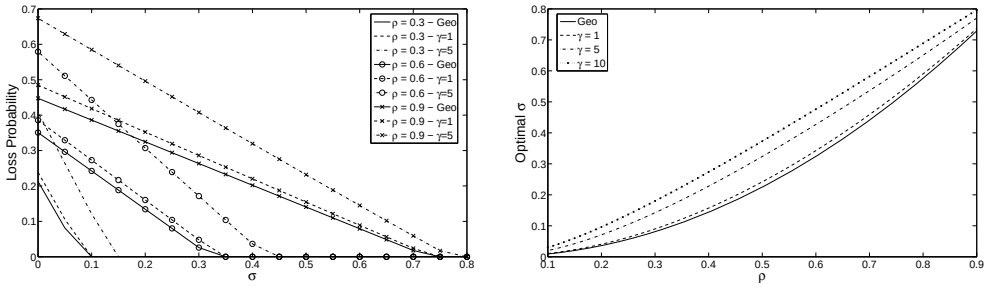

(a) Loss probability. $L \sim \text{Unif}(5, 15)$ (b) Optimal σ. $L \sim \text{Unif}(5, 15)$

Figure 4: Effect of the burstiness

15; and a distribution with two equally likely values, 5 and 15. Figure 3(b) shows that the effect of the packet-size distribution on σ^* is also rather small, a result that holds for a broad range of load values. At this point we must recall that we are considering the homogeneous-traffic case, where all the wavelengths have the same arrival process and packet-size distribution. In this case the results are obtained for $K = 2$, but this parameter has no effect on the results since all the wavelengths are identical and the number of converters is proportional to the total number of wavelengths, which is infinite for any value of K. Therefore, when presenting the results under the assumption of homogeneous traffic we do not need to specify the value of K, as the results are the same for every $K \geq 1$. This also means that when the number of wavelengths tends to infinity and the traffic among the ports is homogeneous, there is no difference between the performance of the centralized and non-centralized architectures.

Although the packet-size distribution seems to have little effect on the performance of the switch, Figure 3 shows a significant difference between the results for geometric and ON-OFF arrivals. We now consider the effect of the arrival process' burstiness in more detail, by means of the ON-OFF process. A simple measure of the burstiness of an arrival process is the ratio between its peak rate and its mean rate [18]. For geometric IATs these two rates are equal and the ratio is one. For the ON-OFF process the peak rate is q (the rate of the geometric IATs during the ON periods), the mean rate is $\frac{q}{\gamma+1}$ and the ratio is $\gamma + 1$. Therefore, increasing the value of γ while keeping the load fixed increases the burstiness of the process, which is expected since the same number of arrivals will occur in shorter time intervals (ON periods), followed by longer silent (OFF) periods. Figure 4(a) depicts the results for geometric and ON-OFF arrivals with various values of γ (one and five) and various loads (0.3, 0.6 and 0.9). We observe that a larger burstiness implies a significantly larger loss probability, and therefore a larger conversion ratio to achieve zero losses. This effect is considerable for any load, but it is particularly relevant for mid loads. The effect of the burstiness on σ^* is shown in Figure 4(b), where the larger absolute effect for mid loads is evident. For instance, for $\rho = 0.9$ the value of σ^* for geometric IATs is around 0.73, while for ON-OFF($\gamma = 5$) arrivals is 0.77. For a load of 0.5, σ^* is 0.22 for geometric IATs and 0.32 for ON-OFF($\gamma = 5$) arrivals. It appears that when the load is high, the main cause of losses is the load and the burstiness only comes in second place.

4.3. Heterogeneous traffic

We now consider heterogeneous traffic conditions, starting with the case of a difference in burstiness among the traffic directed to the output ports. In Figure 5(a) we show the loss proba-

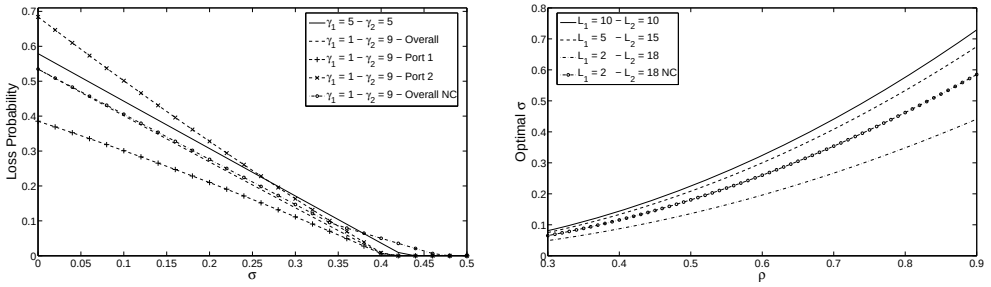

(a) Heterog. burstiness ($\gamma_1 = 1$), $\rho = 0.6$, $L \sim$ U(5, 15)

(b) Heterogeneous packet size. Geo IATs

Figure 5: Effect of heterogeneity - $K = 2$

bility for a switch with two output ports, where the only difference between the traffic directed to these ports is related to the burstiness. We consider two cases: in the first, both output ports have the same ON-OFF arrival process, with parameter $\gamma = 5$; in the second case, the traffic to the first output port is an ON-OFF process with $\gamma = 1$, while for the second port the traffic follows an ON-OFF process with $\gamma = 9$. In the first case, the loss probabilities per port are the same, equal to the overall loss probability. For the second case the ports present a very dissimilar performance, where the larger losses correspond to the port with the more bursty traffic, as expected. When the conversion ratio is highly under-dimensioned the difference between the performance of the ports is very large, as is the overall loss probability. As the conversion ratio increases and gets closer to σ^*, the difference between the loss probabilities of the ports decreases. We also observe that the overall loss probability is smaller for the heterogeneous-traffic than for the homogeneous-traffic case, and there is also a difference in the value of σ^*, favorable to the heterogeneous case. This is particularly relevant since, for the type of arrival processes we are considering, the overall burstiness is the mean of the burstiness among the traffic for all ports. As a result, in both scenarios the overall burstiness is 5, but when the traffic is heterogeneous both the loss probability and σ^* are smaller than in the homogeneous case. In addition, Figure 5(a) also depicts the loss probability for the heterogeneous case when there converters are arranged in a non-centralized fashion (labeled NC). We observe that the overall loss probability is very similar for the centralized and non-centralized cases up to a certain value of the conversion ratio (in this instance 0.35). For larger values of σ the difference increases considerably, with the centralized architecture showing a smaller loss probability and a significantly smaller σ^*.

In Figure 5(b) we also consider heterogeneous traffic conditions but this time the difference is in the packet-size distribution. In this case the load offered to the switch is the same, but there are differences in the mean packet size per port. Figure 5(b) shows the value of σ^* for three cases: in the first case the packet size of both ports is deterministic and equal to 10; in the second the packets for the first port have size equal to 5 while those for the second have size equal to 15; in the last case the packets for the first and second ports are of size 2 and 18, respectively. In all the cases the load is the same for both output ports. We see that the packet size has an important effect on σ^*, particularly for high loads. When the asymmetry in the packet sizes is larger, the conversion requirements decrease, but this is significant only if the asymmetry is sufficiently large. In particular, we observe little difference between the first and the second cases, but a large difference between these two and the third case. The gain in conversion requirements can be partly attributed to the centralized location of the converters, as opposed to having a pool

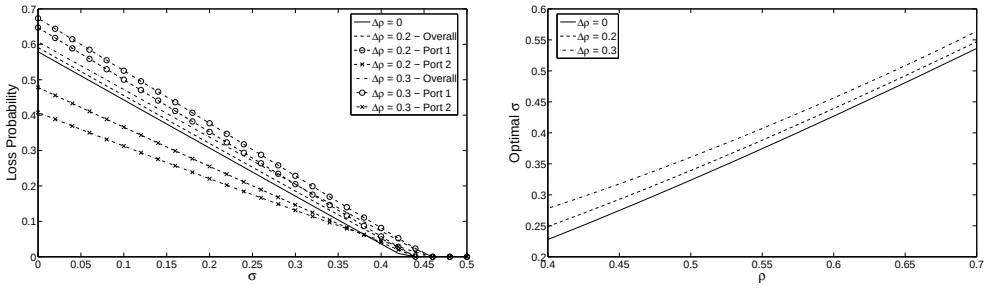

(a) Loss probability. $\rho = 0.6$, $L \sim \text{Unif}(5, 15)$, $\gamma = 5$

(b) Optimal σ^*. $L \sim \text{Unif}(5, 15)$, $\gamma = 5$

Figure 6: Effect of load heterogeneity - $K = 2$

of converters per port. The non-centralized case is illustrated for the third case in Figure 5(b) with the label NC, which shows the conversion requirements to attain a zero loss probability if each port had its own set of converters. As stated above, under the assumption of an infinite number of wavelengths, the number of ports K has no effect on the performance of the switch if the traffic is homogeneous. Therefore, in that case the difference between centralized or non-centralized conversion vanishes as the number of wavelengths becomes large. However, as soon as the heterogeneity in the traffic is considered (either in burstiness or packet-size distribution) we observe an important gain obtained by centralizing the conversion resources. The other reason for having a lower loss probability in the latter two cases is that, as the ports have the same load, the arrival rate for the port with the smaller packet size is larger, and therefore there is a larger proportion of small packets, which means that the mean packet size is also smaller.

As a final scenario we consider heterogeneously loaded ports. We analyze a switch with two ports, where the difference between the load of the first (ρ_1) and second (ρ_2) ports is given by $2\Delta\rho$, with $\Delta\rho = \rho_1 - \rho = \rho - \rho_2$, and ρ is the overall load. Figure 6(a) illustrates the effect of load heterogeneity, under ON-OFF($\gamma = 5$) arrivals, for three values of $\Delta\rho$: 0, 0.2 and 0.3. We observe how the overall loss probability is only slightly affected, while the loss probability per port is significantly different, especially if the conversion ratio is too small compared to σ^*. Also, the effect of $\Delta\rho$ on the minimum conversion ratio to attain a zero loss probability, shown in Figure 6(b), is rather small, even when the asymmetry in the loads is large. In this case, the value of σ^* is identical under non-centralized and centralized conversion, meaning that no gain in conversion resources is obtained by using centralized conversion when the asymmetry in traffic arises from a difference in the load. However, in the non-centralized case the converters would need to be allocated in each port proportionally to the loads. As the load is a very dynamic parameter, a converter allocation based on the mean load per port would result in a combination of periods with many idle converters (lowly loaded) and periods with many losses (highly loaded). Also, an allocation based on the peak load would require a large number of converters per port. In a dynamic-load scenario, the centralized architecture will provide an important gain, as it will be able to combine peak periods for some ports with valley periods for others. In addition, it is highly unlikely to find a real scenario where the only difference in the characteristics among the ports' traffic is the load, while the burstiness and the packet-size distribution are the same for every port. Typically, the difference will be in all these characteristics, and therefore the centralized architecture will provide a better performance and require fewer converters than the non-centralized one, even if the number of wavelengths is large.

References

[1] V. Eramo, M. Listanti, Packet loss in a bufferless optical WDM switch emplying shared tunable wavelength converters, Journal of Lightwave Technology 18 (2000) 1818–1833.

[2] N. Akar, E. Karasan, K. Dogan, Wavelength converter sharing in asynchronous optical packet/burst switching: an exact blocking analysis for markovian arrivals, IEEE J. Sel. Areas Commun. 24 (2006) 69–80.

[3] N. Akar, E. Karasan, G. Muretto, C. Raffaelli, Performance analysis of an optical packet switch employing full/limited range share per node wavelength conversion, in: Proceedings of IEEE Globecom 2007, 2007, pp. 2369–2373.

[4] L. Xu, H. Perros, Performance analysis of an ingress switch in a jumpstart optical burst switching network, Performance Evaluation 64 (2007) 315–346.

[5] S. Yao, B. Mukherjee, S. J. B. Yoo, S. Dixit, All-optical packet-switched networks: a study of contention-resolution schemes in an irregular mesh network with variable-sized packets, in: Proc. of OPTICOMM 2000, 2000, pp. 235–246.

[6] R. Baxter, Exactly solved models in statistical mechanics, Academic Press, 1982.

[7] D. A. Dawson, J. Tang, Y. Zhao, Balancing queues by mean field interaction, Queueing Systems 49 (2005) 335–361.

[8] F. Baccelli, D. R. McDonald, J. Reynier, A mean-field model for multiple tcp connections through a buffer implementing red, Performance Evaluation 49 (2002) 77–97.

[9] C. Graham, P. Robert, Interacting multi-class transmissions in large stochastic networks, Annals of Applied Probability 19 (2009) 2334–2361.

[10] C. Bordenave, D. McDonald, A. Proutiere, Performance of random medium access control, an asymptotic approach, in: SIGMETRICS '08: Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, ACM, New York, NY, USA, 2008, pp. 1–12.

[11] J. Mundinger, J. Le Boudec, Analysis of a reputation system for mobile ad-hoc networks with liars, Performane Evaluation 65 (3-4) (2008) 212–226. doi:http://dx.doi.org/10.1016/j.peva.2007.05.004.

[12] F. Callegati, W. Cerroni, G. Corazza, C. Develder, M. Pickavet, P. Demeester, Scheduling algorithms for a slotted packet switch with either fixed or variable lengths packets, Photonic Network Communications 8 (2004) 163–176.

[13] F. Callegati, W. Cerroni, C. Rafaelli, P. Zaffoni, Wavelength and time domain exploitation for QoS management in optical packet switches, Computer Networks 44 (2004) 569–582.

[14] G. Latouche, V. Ramaswami, Introduction to Matrix Analytic Methods in Stochastic Modeling, SIAM, Philadelphia, PA, 1999.

[15] B. Van Houdt, K. Laevens, J. Lambert, C. Blondia, H. Bruneel, Channel utilization and loss rate in a single-wavelength Fibre Delay Line (FDL) buffer, in: Proceedings of IEEE Globecom 2004, 2004, pp. 1900–1906.

[16] J. Le Boudec, D. McDonald, J. Mundinger, A generic mean field convergence result for systems of interacting objects, in: Proceedings of the 4th QEST, 2007, pp. 3–15.

[17] J. F. Pérez, B. Van Houdt, A mean field model for an optical switch with centralized partial conversion, Tech. rep., University of Antwerp (2010).

[18] T. Robertazzi, Computer Networks and Systems, Springer, 2000.